Two-Sided Matching Agents for Electronic Employment Market Design: Social Welfare Implications

William R. Gates* and Mark E. Nissen

*Contact Author

Graduate School of Business and Public Policy

Naval Postgraduate School

555 Dyer Road, Code GB/Gt, Monterey, CA 93943-5000

Tel: 831-656-2754; fax: 831-656-3407; e-mail: brgates@nps.navy.mil

Two-Sided Matching Agents for Electronic Employment Market Design: Social Welfare Implications

Abstract

Most e-commerce research addresses product and service markets, but the design of equally important employment markets has been largely neglected in the management literature. This article builds upon multidisciplinary research in labor market economics and information systems, and it involves an exploratory experiment to assess the performance of five alternative employment market designs. Using a quasi-price measure for comparison and examining social welfare as a basis for assessing market-design alternatives, we provide novel insight into the balance required between technologically enabled efficiency and economically principled effectiveness of markets. Additionally, this experiment addresses labor market economics in an analytically difficult, public-sector context, yet the market designs and associated technologies are quite general, so the experiment also provides insight into private-sector markets as well. Our results point to a Pareto superior increase in total welfare through market design, which can dramatically increase employee morale and retention, and increase overall labor market efficiency.

1 Electronic Employment Markets

The majority of electronic commerce (e-commerce) research addresses markets and processes associated with the sale of products and services, which is clearly very important for developing a solid understanding in the field. This research draws from the disciplines of economics (e.g., market design, intermediation, pricing) and information systems (e.g., system design, Internet portals, software agents), and it benefits the practitioner as well as adding to our

cumulative knowledge and understanding through research. But in comparison with research addressing product and service markets, an important area of e-commerce has been relatively neglected: employment markets. These latter employment markets offer potential to gain from e-commerce research as much as their product and service counterparts, and they similarly lend themselves to multidisciplinary investigation.

The central task performed by employment markets involves matching people with jobs, and market mechanisms (e.g., balancing supply and demand through adjustable wages) help promote the efficient allocation of job seekers to the jobs making best use of their skills, and vice versa. Although such mechanisms mirror those found in most markets for goods and services, employment markets differ in a fundamental manner: two-sided matching. In matching people with jobs, for instance, the job seeker (i.e., the "good") must want a particular job, in addition to the employer (i.e., the "consumer") wanting to hire this specific person. In markets for goods and services, on the other hand, matching is generally one-sided. For instance, there is no requirement for a particular good (e.g., a hammer) to want a consumer (e.g., a carpenter), in addition to the consumer wanting to purchase the good. This difference complicates the design of employment markets.

The introduction of information technology (IT) to enable electronic employment (e-employment) markets further adds to this design challenge. For instance, market efficiency depends upon good information (i.e., perfect, in theory), but employment markets are very information-intensive (e.g., consider the magnitude of relevant information pertaining to all jobs for software developers and all people with software-development skills). Moreover, most IT tools currently available to employers and job seekers contribute little to mitigating the associated information overload (e.g., when a Web search engine returns 10,000 software-

developer prospects). And although the recent emergence of more powerful IT tools, such as "intelligent" software agents, offers considerable promise in terms of helping boundedly rational people to search and filter through myriad job/employee prospects, important properties (e.g., market clearing) are difficult to ensure through such technologically enabled market designs.

This article builds upon multidisciplinary research in labor market economics and information systems, and it involves an exploratory experiment to assess the performance of five alternative market designs. Employing a quasi-price measure for comparison and examining social welfare as a basis for assessing market-design alternatives, we provide novel insight into the balance required between technologically enabled efficiency and economically principled effectiveness of markets. And through experimentation, we systematically evaluate the relative performance of diverse market-design alternatives (e.g., with and without IT support, distributed software agent automation, centralized algorithms for two-sided matching and optimization).

Additionally, this experiment addresses labor market economics in a difficult, public-sector context, in which the concept *profit* has no bearing, and microeconomic constructs such as *marginal revenue* and *marginal cost* fail to capture many critical employment market considerations (e.g., employer and employee preferences). Yet the market designs and associated technologies are quite general, so the experiment also provides insight into private-sector markets. Indeed, extension to for-profit firms and labor in industry represents a simpler, more-restrictive case of the market designs examined in this investigation.

The balance of the article follows this introduction by highlighting key aspects of employment market economics, after which we outline the kind of software agent technology now emerging from the laboratory to enable a new generation of e-employment markets. We then discuss the two-sided assignment model in some detail, and we lay out the experimental

design. Experimental results are then summarized in turn, and we close the article by discussing

key conclusions and topics for continued research along the lines of this investigation.

2   Labor Market Economics

For relevant background in this section, we highlight key points from what is now

textbook economics (Ehrenberg and Smith 1997). Two modes of matching people with jobs

prevail at present: 1) distributed markets and 2) hierarchical planning. Mirroring most markets

for goods and services, the former approach supports distributed, point-to-point matching

between job seekers and employers. Such, market-based labor markets balance demand and

supply through dynamic wage adjustment. They also allocate labor to its highest valued uses

(i.e., demand efficiency) and to the uses for which each employee is best-suited or most willing

to work (i.e., supply efficiency). Further, compensating wage differentials ensure balance across

industries. In this sense, market-based labor markets maximize social welfare, which is

traditionally measured as the area between the labor demand and supply curves, for the quantity

of labor hired (i.e., the shaded area in Figure 1). Despite these positive properties, however, the

market-based approach is seldom used for employee/job matching *within* modern enterprises,

and current employers have little control over people seeking employment outside the enterprise.
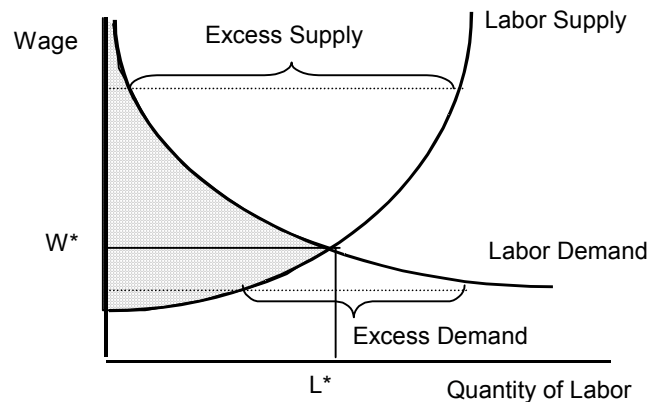
Figure 1. Market-Based Labor Markets

Alternatively, patterned after centrally planned economies (e.g., the former Soviet Union) and command-and-control organizations (e.g., the military, government, large corporations), the latter approach involves administrative procedures to match people with jobs. Hierarchical job assignments rely upon the cognitive processes of centralized, administrative professionals to match individual capabilities and job requirements and to reflect both the job's relative priority and the individual's job preferences. This provides management with tremendous control over when and where specific employees change jobs within the enterprise. But the hierarchical approach is seldom used for employee/job matching *between* modern enterprises, and the centralized matching process lacks many benefits of its market-based counterpart (e.g., market clearing, demand and supply efficiencies).

Thus, distributed market and hierarchical planning approaches each have relative advantages and disadvantages in terms of employer/job seeker matching. In the present research, we investigate an approach that incorporates the best features of both markets and hierarchies in designing employment markets. As suggested in the introduction, two—non-mutually exclusive—approaches are considered: 1) two-sided matching markets and 2) software agent technology.

A third approach, auctions, is not considered here. Auctions are typically used when stable market prices are infeasible because of volatile demand and supply conditions (e.g., stock markets) or limited market participation (e.g., a single buyer with a limited number of sellers, as in government contracting, or a single seller with a limited number of buyers, as in art or antique auctions). Furthermore, there is typically imperfect and asymmetric information so that buyers (sellers) have different views of the item's true value and information about these values is limited (McAfee and McMillan 1987). Of these issues, information anomalies related to

differing valuations of alternative job assignments is the most relevant to our analysis. While the assignment problem could be solved using auctions, this approach is unrealistic in most public sector labor market applications, the primary focus of this research, and likely in private sector internal labor markets as well. These labor markets generally do not have the wage flexibility to allow job seekers or employers to bid for jobs or employees.

3   Two-Sided Matching Markets

A two-sided matching market (Roth and Sotomayor 1990) assigns individuals to jobs when there are several possible employers and job seekers. The matching algorithm balances the employers' and job seekers' preferences, but it can produce assignments that give priority to either employers or job seekers. As such, the algorithm specifically addresses *both* demand and supply efficiency. Two-sided matching algorithms are currently used in assigning medical students to residency programs (Roth and Sotomayor 1990, Roth 1984, Roth and Peranson 1997) and pledges to sororities at some colleges and universities (Mongell and Roth 1991).

The market for medical residents illustrates the two-sided matching system. As U.S. students complete their final year of medical school, they interview for residency positions. Each student interviews with several residency programs, and each program interviews several students. After the interviews, students rank residency programs according to their individual preferences, and programs independently rank students according to their preferences. Students and programs submit their prioritized lists to a central clearinghouse. The clearinghouse compares the lists and assigns students to programs. On a predetermined date, students and residency programs receive their assignments. Each matched student is assigned to one residency program, and each program is assigned students up to the number of available positions.

Unmatched students individually seek unfilled positions; programs with unfilled positions can seek either unmatched U.S. medical students or foreign-trained students.

Participating in this centralized assignment process is voluntary. Residency programs and medical students are free to establish individual agreements, but over 90% of assignments are made through this voluntary, centralized process. To generate and sustain this participation level, the matching process must satisfy a few basic conditions. One of the most important conditions is stability: both students and programs must be at least as happy with their assigned match as with any agreement they could reach outside the centralized process. The outcome is unstable if a student and program both prefer one another to the respective program and student with which they are centrally matched. With unstable matches, the student and program would both choose to forgo the assigned match and form their own agreement. If a student is not matched to his or her highest ranked program, the program must have been assigned students that it ranked more highly (i.e., the program would not reject the assigned match). If a program does not receive its highest ranked students, these students must be matched with programs they rank more highly (i.e., the students would not reject the assigned match). Roth (1991) describes how markets unravel when programs don't meet these requirements.

To summarize the relative advantages and disadvantages, as currently implemented for matching medical students with residency programs, the two-sided matching market incorporates a number of positive features of both hierarchical and market-based labor markets. Most importantly, unlike hierarchical systems, matching markets balance both employers' and job seekers' preferences. This improves both demand and supply efficiencies relative to hierarchical labor markets, and it increases social welfare in labor markets. Two-sided matching markets also are responsive enough to keep pace with active labor markets, yet periodic matching can dampen

the high rates of employee turnover now experienced in high-technology industries. And unlike

market-based systems, two-sided matching markets provide some centralized control through the

clearinghouse. Alternatively, two-sided matching markets lack the automatic, dynamic response

of market-based systems, and the opportunity for side agreements to circumvent the system and

exploit instabilities, if present, can be administratively cumbersome. Further, two-sided matching

markets represent a centralized approach to employment market design, in which market

participants are subject to the rules, timing and potential vagaries of the central-market manager.

4   Software Agent Technology

Work in the area of software agents has been ongoing for some time, and it addresses a

broad array of applications (Bradshaw 1997, Franklin and Graesser 1996, Nissen 2001, Nwana

1996). From this array of applications, software agent technology has been identified as having

particular promise in terms of automation and support in electronic labor markets (Gates and

Nissen 2001). As computational artifacts, agents help overcome human cognitive limitations

(e.g., limited memory and processing speed), and they support rapid search and effective filtering

through huge numbers of available jobs and job seekers.

Further, "intelligent" agents (i.e., possessing some artificial intelligence) can employ

inferential mechanisms (e.g., rules, cases, scripts) to depict and consider diverse individuals'

preferences. In a very large, public-sector enterprise, such as the U.S. Navy (e.g., with several

hundred thousand jobs and sailors, a third of which are subject to personnel rotation in any given

year), in which both employers and job seekers have specific, idiosyncratic needs and

preferences, no other extant information technology offers the automation and support capability

provided by software agents.

Building upon prior research on distributed software agents to automate and support the enterprise supply chain (Nissen 2001), a similar multi-agent system called the Personnel Mall has been developed to enable electronic employment markets. Central to the Personnel Mall's potential in this labor market domain is its ability to represent a multitude and wide variety of different users—on both the demand and supply sides—to quickly find, retrieve and organize large amounts of market information. Its conformance to market and organizational rules, established for a particular enterprise or circumstance, enables this multi-agent system to automate and support commerce in a broad diversity of electronic markets (e.g., including regulation-laden, hierarchical systems). Such ability suggests the Personnel Mall offers good potential to enact, automate and support the kinds of electronic labor markets addressed through this research.

As its name suggests, the Personnel Mall employs a shopping-mall metaphor for employer/job seeker matching. In a mall, "shoppers" (e.g., employers) are not expected to know in advance which "shops" (e.g., job seekers) exist or what "products" (e.g., qualifications) they offer for sale. Neither are the job seekers expected to know *a priori* what jobs are available, nor are the employers expected to know in advance which job seekers may be interested in their job openings.
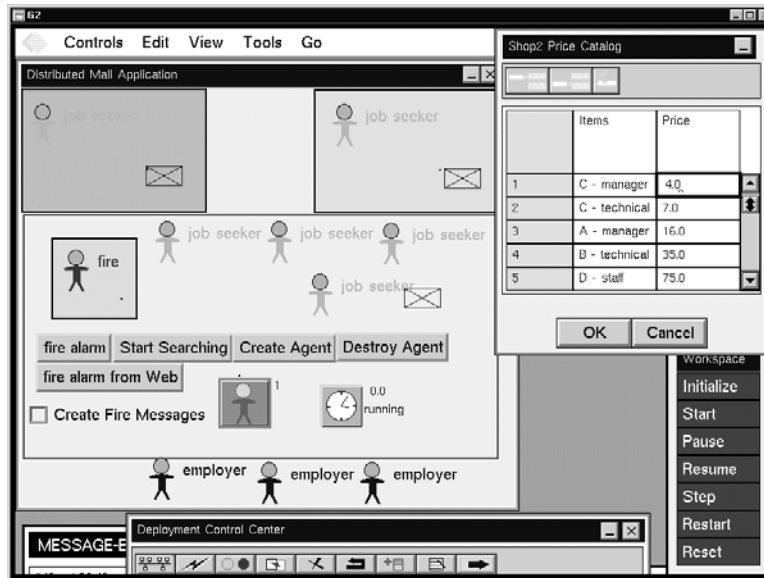
Figure 2. Personnel Mall Screenshot

The Personnel Mall employs a quasi-price system to support comparisons across job seekers. Job seekers use quasi-prices to *quantify* their relative preferences between alternative job assignments. A specific job seeker's input form, as in Figure 2, lists quasi-prices corresponding to the individual's relative job preferences: lower quasi-prices indicate more preferred jobs. This effectively transforms the matching problem from one relying on ordinal scales (e.g., ranked job alternatives) to one quantifying individuals' relative utilities associated with diverse jobs (e.g., using consistent numerical prices, ratings, scores, credits or quasi-pricing instruments).

This is comparable to the manner in which a centralized market (e.g., stock exchange) uses relative prices (e.g., price per share) to compare values across diverse products (e.g., securities), or the manner in which job seekers express their job preferences in market-based labor markets by selecting between jobs based in part on relative wages. This aspect of the Personnel Mall compares favorably with the outcomes of market-based markets discussed above.

The Personnel Mall enables employers' agents to search across various job seekers in the Mall, and job seekers' agents to search across a diversity of participating employers. When viewed from the employer perspective (i.e., employer agents searching across job seekers), for instance, the employer agent uses rules to search for and select the lowest "priced" job seeker(s) for each job. The Mall is always "open," as employers and job seekers can launch agents to search respectively for people and jobs at any time, and each agent will search for the best match available in the Mall at the time. This mirrors the manner in which physical employment markets function (e.g., search is limited to market participants at the time). Where more than one alternative (e.g., job, person) is acceptable, the agent will select the one with the lowest quasi-price. If multiple acceptable alternatives have equivalent quasi-prices, then agents select the first one identified to break quasi-price ties (i.e., first-come-first-matched). From the perspective of the "shopping" agent, this implementation is essentially a one-sided match, which is most appropriate when employers/job seekers view job seekers/jobs as homogenous.

In practice, however, this condition of homogeneity appears to be too restrictive, so we allow employers to have different preferences over job seekers and job seekers to have different preferences over employers. This mimics a more representative employment market, in which employers consider both their own preferences for job seekers and the job seekers' quasi-prices for that employer. Leaving all other Mall mechanics unchanged, agents representing one class (e.g., employers) will select the available agent representing the other class (e.g., job seekers) that provides the highest social welfare of the potential assignments available. Although the mechanics are the same, results may differ depending upon the market mode (i.e., whether employers are "shopping" for job seekers or vice versa). We report results from both market modes in this study.

Overall, the Personnel Mall is subject to many of the same market principles and dynamics found in physical labor markets. As with two-sided matching algorithms, the Personnel Mall appears to offer some improvements over hierarchical planning systems (e.g., explicitly considers employer and job seeker preferences), and it offers advantages only attainable through electronic markets (e.g., provides access to and helps manage search through abundant, market-wide information; automates many search and matching tasks).

However, as with the approaches discussed above, the Mall also suffers from some relative disadvantages. For instance, despite use of a quasi-pricing mechanism, the distributed nature of the Personnel Mall cannot guarantee that labor supply and demand will balance (e.g., markets do not always clear), and the individual-centered behavior of software agents using a fixed search strategy (e.g., lowest quasi-price of participating agents, first-come-first-matched to break quasi-pricing ties) may sacrifice global market efficiency to satisfy the individual preferences of employers and job seekers whose agents are participating in the Mall at any given time.

5   Two-Sided Assignment Model

Testing the performance of different job assignment mechanisms requires measuring the goodness of fit between job seekers and jobs that have been matched with one another. The most relevant metric to measure fit is social welfare (i.e., the area between the labor supply and demand curves). The greater the social welfare, the better the match between employers and job seekers. Although traditional labor supply and demand curves can be employed to delineate supply and demand directly in market-based markets, they do not apply well to hierarchical labor markets (e.g., employers and job seekers are not homogeneous). Hence, as generally employed,

this technique is limited to a special case, representing only a subset of labor markets. Thus, we must develop other labor market performance measures.

Three metrics appear most obvious: 1) the preference rank-order of the assignment for the job seeker, the employer, or both; 2) the utility, or level of satisfaction the job seeker (employer) receives from the assigned job (job seeker); and 3) quasi-prices such as those employed by the Personnel Mall. Preference rank-orders are both measurable and comparable across job seekers and employers, but they don't capture the relative strength of preferences. Utilities capture the relative strength of preferences, but in practice they are not comparable across individual job seekers or employers. Quasi-prices are comparable across individuals, but they must be set, or adjusted, in a manner that reflects the relative strength of preferences.

These empirical difficulties lead us to explore the design space of alternative market approaches using simulated employers and job seekers. This enables us to compute utility values for specific employers and job seekers in a simulated labor market, where utility values can be assigned to ensure comparability across job seekers/employers. Based on discussions with the Economic Science Laboratory at the University of Arizona, and following the work of Ledyard, et al. (1996) and Olson and Porter (1994), we report quasi-prices in place of utilities, which enables matching results to be summed across job seekers/employers. These quasi-prices are equivalent to the prices employers (job seekers) would bid for the associated job seeker (job) under an incentive compatible bidding process.

In the present case, quasi-prices are simply the inverse of the job seeker's or employer's utility values. Hence, higher utilities translate into lower quasi-prices, and vice versa. For job seekers, quasi-prices can be interpreted as reservation wages: the minimum wage required for a job seeker to accept an assignment. Lower reservation wages signal more preferred jobs. For

employers, quasi-prices represent the cost to accomplish a task with a particular job seeker. Lower employer quasi-prices indicate that the employee in question can complete the task more efficiently (e.g., employees all receive the same pay per time period, but better training or personal performance allows some employees to complete the task more quickly).

A Cobb-Douglas utility function is employed to generate utility values in this exploratory research. Cobb-Douglas utility functions are frequently used to represent individual preferences (e.g., employers, job seekers), because they correspond to the characteristics of consumers' tradeoffs and demand observed in both product and labor (input) markets. Most importantly, Cobb-Douglas utility functions with exponents less than one exhibit diminishing marginal utility. Diminishing marginal utility implies that "consumers" (e.g., employers/job seekers) prefer more of a "good" (e.g., favorable job seeker/job attribute) to less, but their satisfaction increases at a decreasing rate as they have more of that "good" (e.g., attribute).

Diminishing marginal utility leads to diminishing marginal rates of substitution (MRS) between two attributes. Substitution implies that employers/job seekers are willing to trade off more of one attribute for less of another; diminishing MRS implies that their willingness to make this trade diminishes the more they trade. As the mix of attributes shifts in favor of one relative to others, extra units of the more plentiful attribute become less valuable, while extra units of the less plentiful attributes become more precious. Consumer preferences typically reflect diminishing MRS (Mas-Colell, Whinston and Green 1995, Ch. 3), and our utility values are computed accordingly. This provides a measure of realism and external validity to our simulated labor markets.

With this well-accepted model, we generate unique utilities for all job seekers and employers. By systematically varying the utility functions' coefficients, two thousand unique

utility values are generated and used to represent the preferences of various simulated job seekers and employers. From the 2000 utility values—1000 for job seekers and 1000 for employers—we randomly select ten job seekers and twelve employers. This creates a scenario with 20% over-demand for labor, which mirrors the challenging conditions experienced in several labor markets (e.g., IT professionals, military personnel, governmental employees). Clearly, using automated approaches such as software agents enables much larger numbers of job seekers and employers to be matched, but matching such large numbers exceeds the capabilities of our human subjects. The job seeker and employer utility functions are summarized in the Appendix.

6 Experimental Design

The research design involves an exploratory experiment to assess five alternative designs for employment markets. Given the lack of research to date focusing specifically on *electronic* employment market design, however, the extant literature has not yet developed to the level at which informative market-design research hypotheses can be developed and tested. This is the basis for pursuing exploratory experimentation. Yet experimentation represents a powerful research method, one that empowers the investigators with substantial control over research conditions (e.g., subjects, tasks, environment, variables). Thus, we expect to gain a wealth of new knowledge pertaining to the design of electronic labor markets.

To assess a variety of market designs, the experiment focuses specifically on the task of matching job seekers with job openings, for the quality of such matching determines the utilities of both employers and job seekers. And using our quasi-pricing scheme, higher utilities translate into greater social welfare from a particular labor market design. Additionally, we examine five different experimental conditions, each of which is associated with a unique market design

approach, and we describe the human subjects participating in the experiment. The experimental procedure is outlined subsequently, and instead of formal research hypotheses, we articulate five expectations for results of the experiment. The experimental context is established to match that of a hierarchical labor market.

6.1   Experimental Conditions

Five experimental conditions are examined: 1) unassisted, 2) assisted, 3) Personnel Mall, 4) two-sided matching algorithm, and 5) optimization. The first two experimental conditions involve people performing the matching task, whereas the other three automate this task, albeit through alternative mechanisms. All conditions involve the identical set of twelve job openings and ten job seekers participating in the market. As the name implies, the unassisted condition is used to assess the performance of people performing the matching task with no technological or algorithmic support. This represents a control group, which is used for comparison with the other conditions, and this experimental condition corresponds with the manner in which most (hierarchical) job matching is currently accomplished in practice. Alternatively, subjects in the assisted condition use a product called Logical Decisions for Windows (Logical Decisions 1993) to assist them with the matching task. This tool enables multi-attribute-utility analysis, and it supports both graphical and numerical analysis of complex decision problems (e.g., involving many alternatives, attributes, performance levels). All subjects are taught to use this tool proficiently before the experiment.

The third experimental condition automates the matching task through the Personnel Mall, which uses software agents to represent both employers and job seekers, and quasi-prices (i.e., inverse utilities) to represent employer and job seeker preferences. Using these quasi-prices, the Personnel Mall is run twice to reflect respective employer and job seeker biases.   In

the former run, employer agents search for and match to job seekers by comparing the employer's value for a job seeker to that job seeker's quasi-price for that employer. In essence, this comparison measures the net social welfare generated by prospective employer/job seeker matches. Assignments are made in (randomly arranged) employer order; that is, Employer Agent 1 first makes its preferred match, then Employer Agent 2 makes its preferred match, and so forth until all ten job seekers are assigned. Alternatively, in the latter run, job seeker agents search for and match employers based on the job seeker's preferences relative to the employers' quasi-prices, and assignments are made in (randomly arranged) job seeker order; that is, Job seeker Agent 1 first makes its preferred match, then Job Seeker Agent 2 makes its preferred match, and so forth until all job seekers are assigned. This assignment methodology mimics a first-come first-served approach, which characterizes many public- and private-sector, hierarchical job-assignment processes. The employer and job seeker biased processes where both run two hundred times, with randomized arrival orders for employers and job seekers, respectively.

The fourth experimental condition automates the matching task through a two-sided matching algorithm, which is set up to simultaneously consider the preferences of all employers and job seekers. It also establishes a set of matches that balances both employer and job seeker preferences in order to ensure stability. Like the Personnel Mall, the matching algorithm is run twice: first, the algorithm is used with an employer bias; the second run is conducted with a job seeker bias. In this specific experiment, the matching results from both runs are identical. This similarity in matching results is common in practice and characterizes labor markets in which employers and job seekers maintain realistic expectations for the kinds of job seekers and jobs, respectively, that they list in their preferences (Roth 1999).

The fifth experimental condition automates the matching task through an optimization algorithm, which explicitly seeks to minimize average quasi-price across the entire set of employers, job seekers, or both. The optimization algorithm is run three times: once to minimize the job seekers' average quasi-price, once to minimize the employers' average quasi-price, and once to minimize the average of job seekers' and employers' combined quasi-prices. Average rather than total quasi-prices are used to maintain a consistent scale across experimental conditions.

6.2   Human Subjects

The experiments to measure manual task performance are conducted using 54 graduate students in the management curriculum of a U.S. university. Ten subjects participate in the same experiment a second time, which is useful to examine any learning effects on matching performance. Additionally, six employment-matching professionals also participate in the experiment, which is useful to examine any performance differences between amateurs (e.g., graduate students) and professionals in this task. Further, twelve students in a different MBA program also participate in the experiment, which is useful to examine any performance differences associated with a particular class of students. A similar sample of 22 management students participated in the assisted trials.

All of the subjects have an undergraduate degree, 5 – 15 years' work experience, and current middle-management positions with large employers. This group is quite representative of people that currently participate in employment markets—hierarchical as well as market-based—both as job seekers and employers. Therefore, the results of this experiment are expected to generalize well to a broad range of employment markets in industry.

6.3   Experimental Procedure

Procedurally, the participants receive a short introduction to the matching process and the Personnel Mall, to familiarize them with agent technology and its potential application to the experimental task. Following this introduction, participants are told about the rationale for and nature of the experiment, and the experimental procedures are explained. Each participant receives an individualized package of experimental materials (e.g., instructions, listing of job openings, job-assignment forms), and each item is described. Participants are encouraged to ask questions throughout this introduction, and an extended period for questions and answers is included before the experiment. The student subjects participate in the study as a part of a graduate school class. They are encouraged to perform well, as results are incorporated into their course-grading scheme.

To make the job assignments, participants are instructed to complete a simple form. The form lists each job and requires the subject to fill in the job seeker assigned. Participants are given an opportunity to ask questions pertaining to matching, the experiment or anything else that may be of concern or confusion. The participants' reactions and comments are collected after the experiment to provide feedback for future experiments.

6.4   Experimental Expectations

Despite our use of exploratory experimentation in this study (e.g., absence of formal research hypotheses), we know enough about the different kinds of technologies and algorithmic capabilities associated with our experimental conditions to develop a set of expectations for the experiment. Specifically, considering the five experimental conditions, we would expect that the unassisted control group is likely to have the poorest matching performance (e.g., highest average quasi-prices and lowest social welfare). Next, we would anticipate assignment

19

performance to increase for the assisted group, because the decision-support tool enables subjects to systematically capture preferences and compare alternatives. However, the assignments in this experimental condition remain bounded by the cognitive limitations and rationality of human subjects. Alternatively, performance for the Personnel Mall has good potential to exceed that of both human-subject groups, for it explicitly considers quasi-prices in its search for employer/job seeker matches. However, the first-come first-served assignment algorithm used in this experiment limits that potential.

Assignment performance is likely to be the best for the two-sided matching and optimization approaches, because these approaches incorporate the batch-processing capability currently lacking in the Personnel Mall. Between these two algorithmic approaches, two-sided matching may not generate higher social welfare than the optimization approach, but the matches produced will be stable and hence inherently feasible in practice. Conversely, optimization is expected to generate the highest social welfare of all approaches, but it may not produce stable matches.

7   Results

Results are summarized in Table 1. The unassisted condition reflects pooled data for all 80 subjects. Pooling the data establishes a sizeable control group for comparison with performance in the other experimental conditions. Pooling is justified in this case, because performance differences between the various human-subject groups are not statistically significant. Similarly, results in the unassisted condition reflect pooled data.

Notice that subjects in the control group perform better (i.e., produce lower quasi-prices) for employers than they do for job seekers, and this difference is statistically significant at the 99% level. Specifically, the aggregate quasi-price for job seekers is $3.92, whereas subjects achieve a better aggregate of $3.48 for employers. Apparently, subjects have some bias toward

serving employer interests above those of job seekers. The job matching professionals surveyed

in our early pilot tests indicate that such a bias toward employers is very common in practice.

Hence, these results are consistent with professional practice in employment markets and

expected to generalize well beyond this experiment. Total quasi-price (i.e., $7.40) is calculated

as the sum of quasi-prices pertaining to the job seeker and employer values.

Table 1. Experimental Results

Experimental Condition	Aggregate Job Seeker Quasi-Price	Aggregate Employer Quasi-Price	Total Quasi-Price
Unassisted	$3.92	$3.48	$7.40
Assisted	$3.87	$3.45	$7.32
P-Mall – Employer	$3.87	$3.37	$7.24
P-Mall – Job Seeker	$3.11***	$2.99***	$6.10***
Matching Algorithm (both runs)	$3.00***	$3.16***	$6.16***
Optimization – Employer	$3.91	$2.52***	$6.43***
Optimization – Job Seeker	$2.68***	$4.04***	$6.71***
Optimization – Combined	$2.94***	$2.76***	$5.70***

*** Significant at 99%

Subjects in the assisted condition maintain this employer bias. The aggregate quasi-price

for job seekers is $3.87, whereas employers are scored at $3.45, though this difference is only

significant at the 90% level. Furthermore, as expected, performance for this group is slightly

better than exhibited by the control group, for both job seekers and employers, but the

differences are not statistically significant. The LDW tool used in this assisted group helps to

make relative preferences and corresponding utility values explicit, so subjects may obtain better

insight into the tradeoffs between job seeker preferences, but there is no support to

systematically make these tradeoffs simultaneously for all market participants.

As explained above, the Personnel Mall is run for two hundred trials with both an

employer and job seeker bias. In employer-biased runs, the results are slightly better than the

control case, for both job seekers and employers, but the difference is not statistically significant

for job seekers and only significant at the 80 percent level for employers. In contrast, the job seeker biased runs show improvement for both job seekers and employers that is statistically significant at the 99 percent level. Comparing these two variants illustrates the vagaries of a first-come-first-served assignment process in which there are more employers than job seekers.

When there are more employers than job seekers, a first-come-first-served employer biased process will always satisfy the first ten employers that enter the system; the last two jobs remain unfilled and are never considered by any job seekers. If the last two employers are highly desired by at least a few well-suited job seekers, overall assignment performance will be adversely affected. This inefficiency is less problematic for the job seeker biased variant. In this case, all ten job seekers are assigned jobs, regardless of their arrival order. Furthermore, each job seeker considers all available jobs. No employers are precluded from consideration with the job seeker bias. Thus, the job seeker biased variant should be expected to generate better overall job assignments, for both job seekers and employers, than the employer-biased variant.

Considering further the relatively more efficient job seeker bias variant, assignment performance improved for both job seekers and employers, but job seeker assignment performance improved relatively more. These trends are consistent with the two automation-enabled performance effects that we anticipated: automation improves overall efficiency and eliminates employer bias. This indicates that the employment domain offers good potential for automation through the kinds of software agents examined in this study, though improvements in performance depend critically on the assignment algorithm embedded in the software agents. In part, this dependence motivates the two-sided matching and optimization algorithms analyzed below. Either of these algorithms could be embedded in future versions of the Personnel Mall.

The matching algorithm produces lower and statistically-significant job seeker quasi-prices than the control group ($3.00 versus $3.92), and again this improvement for job seekers does not come at the expense of employers; employer quasi-prices under two-sided matching ($3.16) are significantly lower than those of the control group ($3.48). In terms of Pareto optimality, the matching algorithm offers an attractive overall approach to job seeker/employer matching, even though total quasi-price ($6.16)—reflecting total utility in the market—is higher by a statistically significant margin than obtained by the Personnel Mall in the job seeker biased variant ($6.10).

Finally, the results from the three optimization runs are as expected. In all three runs, optimization generates the lowest aggregate quasi-prices, *with respect to each particular bias* (i.e., employer, job seeker, combined), but none of the optimization runs generates the lowest quasi-prices *across all three biases*. Specifically, the employer optimization run provides the lowest employer quasi-prices (i.e., $2.52), but these prices come at the job seekers' expense (i.e., $3.91). The job seeker optimization offers the lowest job seeker quasi-prices (i.e., $2.68), but they come at the employers' expense (i.e., $4.04). The run optimizing combined employer/job seeker utility offers the lowest total quasi-price (i.e., $5.70), but not the lowest possible quasi-prices for either of the constituent groups. As with the matching algorithm, employers and job seekers are better off as a combined group under this result, and these improvements are statistically significant when compared to those of the control group.

Two drawbacks to optimization, as opposed to two-sided matching, involve stability and data requirements. Considering stability, optimized solutions are not necessarily stable. In particular, there are 15 blocking pairs in the job seeker optimization assignments; that is, there are 15 potential employer/job seeker matches where both parties would prefer to be matched to

one another than to their assigned partner in the job seeker optimization result. If any of these pairs identifies one another (e.g., through phone inquiries or trade shows), the job seeker optimization result will unravel. Similarly, there are 12 blocking pairs in the employer optimization result and six in the combined optimization run. Thus, although these results are optimal, they are unstable. Hence, they may not be feasible to sustain in practice.

Considering data requirements, optimization requires preference data for job seekers and employers (e.g., utility values, quasi-prices, etc.), and preference specifications must be comparable across job seekers or employers; two-sided matching only requires rank-ordered lists. It is clearly feasible to obtain rank-ordered preference lists from job seekers and employers, as has been done in the U.S. medical residency market for over four decades; it is far from obvious that we can gather accurate, comparable preference data. Thus, comparing two-sided matching and optimization involves trade-offs between efficiency, stability and data requirements; optimization has the potential to provide more efficient matches, but the solutions are likely unstable and entail much more demanding data requirements.

Table 2 summarizes job seeker, employer and total social welfare for each of the experimental conditions; the percentage increases from the control group are shown in parentheses for each instance where the differences from the control group are statistically significant at the 99% level. The pattern of these results is the same as described above. Most notably, the increase in total social welfare is 9.8% for the matching algorithm and 13.5% for combined optimization. If these comparative results hold as we move from an experimental setting to application in the public and private sectors, the absolute increase in welfare could be dramatic. Results summarized in the table further suggest increases in social welfare for job

seekers are even more impressive (e.g., over 15% with two-sided matching and up to 20% with optimization).

Table 2. Social Welfare per Assignment (change from unassisted control group)

Experimental Condition	Aggregate Job Seeker Social Welfare		Aggregate Employer Social Welfare		Total Social Welfare	
Unassisted	$6.08		$6.52		$12.60	
Assisted	$6.13		$6.55		$12.68	
P-Mall – Employer	$6.13		$6.63		$12.76	
P-Mall – Job Seeker	$6.89***	13.4%	$7.01***	7.5%	$13.90***	10.4%
Matching Algorithm	$7.00***	15.2%	$6.84***	4.8%	$13.84***	9.8%
Optimization – Employer	$6.09		$7.48***	14.7%	$13.57***	7.7%
Optimization – Job Seeker	$7.32***	20.5%	$5.96***	-8.5%	$13.29***	5.5%
Optimization – Combined	$7.06***	16.2%	$7.24***	11.0%	$14.30***	13.5%

*** Significant at 99%

According to the U.S. Department of Defense (DoD) FY2002 budget, the U.S. Navy, which is currently moving from a hierarchical to a more automated assignment process, will spend over $20 Billion on labor in FY2002; the DoD will spend over $82 Billion (U.S. 2002). A shift away from unassisted matching (e.g., that increases labor welfare by up to 20% and total welfare by 10 - 14%) would represent a dramatic increase in employee morale and retention, and a significant improvement in the efficiency with which such public-sector enterprises use taxpayer dollars. Moreover, to the extent that the kinds of electronic employment market designs discussed in this article can be further extended to the private sector, the potential economic impact in terms of social welfare is staggering.

8   Conclusion

Most e-commerce research addresses product and service markets, but the design of equally important employment markets has been largely neglected in the management literature. This article builds upon multidisciplinary research in labor market economics and information

systems, and it involves an exploratory experiment to assess the performance of five alternative employment market designs. Employing a quasi-price measure for comparison and examining social welfare as a basis for assessing market-design alternatives, we provide novel insight into the balance required between technologically enabled efficiency and economically principled effectiveness of markets. Additionally, this experiment addresses labor market economics in an analytically difficult, public-sector context, yet the market designs and associated technologies are quite general, so the experiment also provides insight into private-sector markets as well. Our results point to a Pareto superior increase in total welfare through market design, which can dramatically increase employee morale and retention, and increase overall labor market efficiency.

Further, our results indicate that both technological and algorithmic support can have a profound effect on matching quality in employment markets, and these results point toward developing hybrid market solutions (e.g., software agents employing two-sided matching or optimization) as a promising step for future research along the lines of this investigation. Even so, our experiments involve a very small, simulated labor market and simplistic preference functions. Promising future research could also examine large labor markets where preferences are diverse and based on a complicated set of job seeker and employer characteristics, for instance, and research to analyze the behavior of diverse physical and electronic employment markets outside of the simulator is further warranted.

In terms of generalizability, our results focus on the kinds of hierarchical labor markets that are most common within large organizations, but nothing in our approach or results would seem to limit generalization across a broad range of such organizations (e.g., large and small, from the public and private sector). Alternatively, there is some question as to how well

techniques such as the matching algorithm (e.g., due to the need for ordinal rankings) and Personnel Mall (e.g., due to the use of quasi-prices instead of market wage rates) would perform in the kinds of *external* labor markets that are common across organizations and industries. Although we can argue that our approaches and technologies *should* be able to generalize, it remains for a future experiment to demonstrate that they *can* or *do.*

9   References

Bradshaw, J.M. (Ed.), *Software Agents*. AAAI Press: Menlo Park, CA (1997).

Ehrenberg, R.G. and Smith, R.S., *Modern Labor Economics: Theory and Public Policy*. Reading, MA: Addison-Wesley, 1997.

Franklin, S. and Graesser, A. "Is It an Agent or Just a Program? A Taxonomy for Autonomous Agents," in *Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages* Springer-Verlag: New York, NY (1996).

Gates, W.R. and Nissen, M.E., "Designing Agent-based Electronic Employment Markets," *Electronic Commerce Research Journal*, Special Issue on Theory and Application of Electronic Market Design (2001), pp. 239-263.

Ledyard, J.O., Noussair, C., and Porter, D., "The Allocation of a Shared Resource Within an Organization," *Economic Design*, 2 (1996), pp. 163-92.

Logical Decisions, *Logical Decisions: Multi-Measure Decision Analysis Software* Logical Decisions: Golden, CO (1993).

Mas-Colell, A., Whinston, M.D. and Green, J.R., *Microeconomic Theory*, New York, NY: Oxford University Press, 1995.

MacAfee, R.P. and McMillan, J., "Auctions and Bidding," *Journal of Economic Literature*, 25 (1987), pp. 699-738.

Mongell, S.J. and Roth, A.E., "Sorority Rush as a Two-Sided Matching Mechanism," *American Economic Review*, 81 (1991), pp. 441-464.

Nissen, M.E., "Agent-Based Supply Chain Integration," *Journal of Special Topics in: Information Technology & Management* 2:3 Special Issue, Electronic Commerce in Procurement and the Supply Chain (2001), pp. 289-312.

Nwana, H.S., "Software Agents: An Overview," *Knowledge Engineering Review* 11:3 (1996), pp. 205-244.

Olson, M. and Porter, D., "An Experimental Examination into the Design of Decentralized Methods to Solve the Assignment Problem With and Without Money," *Economic Theory*, 4 (1994), pp. 11-40.

Roth, A.E., "A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians in the U.K.," *American Economic Review*, 81 (1991), pp. 415-40.

_____, "The Evolution of the Labor Market for Medical Interns and residents: A Case Study in Game Theory," *Journal of Political Economy*, 92 (1984), pp. 991-1016.

Roth, A.E. and Peranson, E., "The Effects of Change in the NRMP Matching Algorithm," *The Journal of the American Medical Association*, 278 (1997), pp. 729-33.

_____, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, 89 (1999), pp. 748-80.

Roth, A.E. and Sotomayor, M.A.O., *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge, UK: Cambridge University Press, 1990.

U.S. Department of Defense, http://www.dtic.mil/comptroller/fy2002budget/, 2002.

10   Appendix – Utilities

The job seeker's individual utility is given by the following Cobb-Douglas utility function:

$$U_S = P^p * L^l$$

Where:

U_S = Utility of job seeker

P^p = Job seeker's derived utility from the job's promotion prospects

P = Job's index of promotion potential (job visibility)

p = Percent change in job seeker's utility from one percent change in promotion potential

L^l = Job seeker's derived utility from fit between job seeker's preferred location (PL) and job location

L = Job's index of physical location relative to preferred location

l = Percent change in job seeker's utility from 1% change in job location preference

$p + l = 1$

The parameter for promotion potential (P) ranges from 1 to 5, with 5 being the highest and 1 being the lowest chance of promotion through a particular job. Although the 1 – 5 scale and corresponding numerical assignments are somewhat arbitrary, they are logical (e.g., higher values are assigned to better matches) and consistently applied (e.g., every match is scored in the same manner, and every parameter has the same scale). This is sufficient for generating a field of well-behaved utility values for the experiment.

Similarly, job seekers specify their preferred location, and employers specify the jobs' physical locations. If there is a match between locations, the Location Index (L) will be assigned the value of 5; otherwise L will be less than 5 (e.g., minimum value of 1), with the difference in value reflecting the distance between preferred and physical locations. The job seeker will thus derive a higher level of utility for a successful location match, and the closer a job's location is to a job seeker's preference, the higher the value.

Total utility values for each job seeker range from five (i.e., perfect job fit) to one (i.e., worst possible job fit). Correspondingly, individual job seeker quasi-prices range from $0.20 (i.e., reservation wage required to accept the perfect job in this industry) to $1.00 (i.e., the reservation wage required to accept the worst possible job in this industry). With ten job seekers, the aggregate quasi-price for job seekers ranges from $2.00 to $10.00. This analysis also uses the difference between the job seeker's quasi-price and the maximum reservation wage (i.e., $1.00) to measure the fit between the assigned job and the job seeker's preferences (i.e., the job seekers' increase in welfare from receiving a more preferred job). This implies that all job seekers are paid $1.00, but they would be willing to work for as little as $0.20 if they received their most preferred job. The aggregate social welfare for ten job seekers ranges from zero to $8.00.

The employer's individual utility is given by the following Cobb-Douglas utility function:

$$U_C = S^S * E^e$$

Where:

U_C = Utility of employer

S^s = Employer's derived utility from getting job seeker with skill set (S)

S = Employer's index of employee's skill set relative to preferred skill set

s = Percent change in employer's utility from one percent change in skill set index

E^e = Employer's derived utility from getting job seeker with higher performance

index (E)

e = Percent change in employer's utility from one percent change in performance

index

$s + e = 1$

As above, a variety of employer utilities are generated to reflect a tradeoff between the skill set (S) of a job seeker and the job seeker's expected performance level (measured by an evaluation index – E). If a job seeker's skill set matches that desired by the employer exactly, a value of 5 is assigned; otherwise S will be less than 5 (e.g., minimum value of 1), with the difference in value reflecting the distance between preferred and manifest skills. The evaluation index is developed in a similar manner. High-performance job seekers (e.g., with very impressive resumes and references) are scored with a value of 5, and lower-performance job seekers receive proportionately lower scores (e.g., minimum value of 1).

As with job seekers, total utility values for each employer range from five (i.e., perfect job seeker fit) to one (i.e., worst possible job seeker fit). Correspondingly, individual employer quasi-prices range from $0.20 (i.e., the cost to complete a task with the most qualified job seeker) to $1.00 (i.e., the cost to complete the task with the least qualified job seeker). With ten jobs filled, the aggregate quasi-price for employers ranges from $2.00 to $10.00. This analysis uses the difference between the employer's quasi-price and $1.00 to measure the fit between the assigned job and the employer's preferences (i.e., the employer's increase in welfare from receiving a more qualified job seeker). This implies that all job seekers are paid $1.00, but more qualified job seekers can complete the task in twenty percent of the time required by the least

qualified job seekers (i.e., for $0.20). The aggregate social welfare for ten jobs filled ranges

from zero to $8.00.